REPORT DOCUMENTATION PAGE						*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*

4. TITLE AND SUBTITLE
BIOINSPIRED CONCEPTS: UNIFIED THEORY FOR COMPLEX BIOLOGICAL AND ENGINEERING SYSTEMS

5a. CONTRACT NUMBER
FA9550-05-1-0032

5b. GRANT NUMBER
AFOSR

5c. PROGRAM ELEMENT NUMBER

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHOR(S)
Dr. John Doyle

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Caltech
1200 East California Blvd
Pasadena CA 91125

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AFOSR/NL
875 Randolph
Arlington VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-SR-AR-TR-08-0440

12. DISTRIBUTION/AVAILABILITY STATEMENT

13. SUPPLEMENTARY NOTES
Distribution A: Approved for Public Release

14. ABSTRACT
The overall long-term objective of our research is to develop a theoretical foundation for understanding highly evolved and organized complexity, emphasizing the role of network architecture, with application to bio, eco, tech, and social networks, and with strong connections to real-world data. Network-enabled technology can provide unprecedented levels of performance, efficiency, sustainability, and robustness in all domains of science and engineering, and in many areas has already done so. At its best, engineering, biology, and medicine are being transformed by sophisticated platforms to rapidly design/demo/debug/deploy systems of bewildering complexity that are nevertheless typically robust and functional.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER *(Include area code)*

Final Report FA9550-05-1-0032

BIOINSPIRED CONCEPTS: UNIFIED THEORY FOR COMPLEX BIOLOGICAL AND ENGINEERING SYSTEMS

John C. Doyle
Control and Dynamical Systems
California Institute of Technology

The overall long-term objective of our research is to develop a theoretical foundation for understanding highly evolved and organized complexity, emphasizing the role of network architecture, with application to bio, eco, tech, and social networks, and with strong connections to real-world data. Network-enabled technology can provide unprecedented levels of performance, efficiency, sustainability, and robustness in all domains of science and engineering, and in many areas has already done so. At its best, engineering, biology, and medicine are being transformed by sophisticated platforms to rapidly design/demo/debug/deploy systems of bewildering complexity that are nevertheless typically robust and functional. Yet the problems of rare but catastrophic real-world failures and "unintended consequences" are, if anything, becoming worse. This "robust yet fragile" (RYF) feature of complex systems is ubiquitous, and must be faced by any methodology that hopes to be scalable and evolvable with systematic and formal verification approaches.

The development of a theory of architecture will provide essential new insights and tools for understanding the robustness and evolvability as well as fragilities of complex networks. Deeply understanding RYF complexity necessarily means understanding network architecture, an essential and widely discussed but poorly formalized concept. Here, architecture involves the most universal, high-level, and persistent elements of organization, usually defined in terms of protocols—the rules that facilitate interaction between the diverse components of a network. We have compared successful architectures from biology, ecology, and technology, and will expand this dialog across domains. We have already identified a variety of common characteristic organizational structures, including layering of control, with protocols structured as hourglasses and bowties to facilitate robustness and evolvability. This natural history or comparative physiology of architecture motivates the mathematical theory necessary to deepen our understanding of these striking observations, and our preliminary results are encouraging.

Mathematical foundations build on unifying formal analysis methods in cyber/digital, physical/analog, and networked models using robust control theory, dynamical systems, information theory, game theory, numerical analysis, operator theory, real algebraic geometry, computational complexity theory, duality and optimization, and semi-definite programming, motivating new interactions between these diverse areas. Specific applications include robustness analysis of diverse complex control systems in biology and technology, the performance of Internet protocols and their extensions to wireless and ad hoc networks, router topologies and web layout, to wildfire ecologies, to biological signal transduction, stress response, metabolic control, and disease dynamics.

Publications (from 2006)

S. Prajna. Barrier Certificates for Nonlinear Model Validation. Automatica 42(1):117-126. 2006.

D. Alderson, W. Willinger, L. Li, and J. Doyle. The Role of Optimization in Understanding and Modeling Internet Topology. In Telecommunications Planning: Innovations in Pricing, Network Design and Management. S. Raghavan and G. Anandlingham, eds. Springer, 2006.

Towards a Theory of Scale-Free Graphs: Definition, Properties, and Implications, L. Li, D. Alderson, J. C. Doyle and W. Willinger, Internet Mathematics, 2(4), p 431-523, March 2006.

H. El-Samad, A. Papachristodoulou, S. Prajna, J. C. Doyle, and M. Khammash. Advanced Methods and Algorithms for Biological Networks Analysis. Proceedings of the IEEE, vol. 90, number 4, pages 832–853.

H. El-Samad, M. Fazel, X. Liu, A. Papachristodoulou and S. Prajna, Stochastic Reachability Analysis in Complex Biological Networks. Proceedings of the American Control Conference (ACC), June 2006.

L. Li, D. Alderson, J. Doyle, and W. Willinger. Towards a Theory of Scale-Free Graphs: Definition, Properties, and Implications. Internet Mathematics. 2006.

L. Chen, S. H. Low, M. Chiang and J. C. Doyle. Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks. IEEE Infocom, Barcelona, Spain, April 2006.

L. Chen, S. H. Low and J. C. Doyle, On Asymptotic Optimality of Dual Scheduling Algorithm in A Generalized Switch, Proceedings of WiOpt, April 2006.

H. Kurata, H. El-Samad, R. Iwasaki, H. Ohtake, J. C. Doyle, I. Grigorova, C. A. Gross and M. Khammash, Module-Based Analysis of Robustness Tradeoffs in the Heat Shock Response System. PLOS Computational Biology, vol. 2, issue 7, 2006.

L. Chen, S. H. Low and J. C. Doyle, Dual Scheduling Algorithm in a Generalized Switch: Asymptotic Optimality and Throughput Optimality. In High-Performance Packet Switching Architectures, Itamar Elhanany and Mounir Hamdi (eds.), Springer, August 2006.

M. Fazel, D. Gayme and Mung Chiang, Transient Analysis for Wireless Power Control. Proceedings of the IEEE Globecomm, November 2006.

D. Gayme, M. Fazel and J. C. Doyle, Complexity in Automation of SOS Proofs: An Illustrative Example. Proceedings of the IEEE CDC, December 2006.

A. D. Ames and S. Sastry, Hybrid Geometric Reduction of Hybrid Systems. Proceedings of IEEE CDC, December 2006.

P. Tabuada, A. D. Ames, A. Julius and G. Pappas, Approximate Reduction of Dynamical Systems, Proceedings of the IEEE CDC, December 2006.

M. Lobo, M. Fazel and S. Boyd, Portfolio Optimization with Linear and Fixed Transaction Costs. Annals of Operations Research, vol. 152, number 1, pages 376–394, 2007.

L. Chen, T. Ho, S. H. Low, M. Chiang and J. C. Doyle, Optimization Based Rate Control for Multicast with Network Coding, Proceedings of IEEE Infocom, May 2007.

T. Cui, T. Ho and L. Chen, On Distributed Distortion Optimization for Correlated Sources with Network Coding, Proceedings of IEEE ISIT, June 2007.

A. D. Ames, R. D. Gregg, E. D.B. Wendel and S. Sastry, On the Geometric Reduction of Controlled Three-Dimensional Bipedal Robotic Walkers. Lecture Notes in Control and Information Sciences, F. Bullo and K. Fujimoto (eds.), Springer, 2007.

S. Ganesh, A. D. Ames and R. Bajcsy, Composition of Dynamical Systems for Estimation of Human Body Dynamics. In Hybrid Systems: Computation and Control, volume 4416, pages 702–705. Springer, 2007.

A. D. Ames, Homotopy Meaningful Hybrid Model Structures. In Topology and Robotics, M. Burger and M. Farber and R. Ghrist and D. Koditschek (eds.), AMS Contemporary Mathematics Series, 2007.

M. Chiang, S. H. Low, A. R. Calderbank and J. C. Doyle, Layering as Optimization Decomposition. IEEE Proceedings, volume 95, issue 1, 2007.

N. C. Martins, M. A. Dahleh and J. C. Doyle, Fundamental Limitations of Disturbance Attenuation in the Presence of Side Information. IEEE Trans. Automatic Control, vol. 52, no. 1, pages: 56-66 . January 2007.

J. C. Doyle and Marie Csete, Rules of engagement, Nature, volume 446, issue 7138, page 860, April 2007.

On the Diversity of Graph with Highly Variable Connectivity, D. Alderson and L. Li, Physical Review E, vol.75, Issue 4, April 2007

A. D. Ames and R. D. Gregg, Stably Extending Two-Dimensional Bipedal Walking to Three Dimensions. Proceedings of IEEE ACC, July 2007.

H. Sandberg, J.-C. Delvenne and J. C. Doyle, The Statistical Mechanics of Fluctuation-Dissipation and Measurement Back Action. Proceedings of IEEE ACC, July 2007.

J.-C. Delvenne and H. Sandberg and J. C. Doyle, Thermodynamics of linear systems. Proceedings of the European Control Conference, 2007.

H. Sandberg, J.-C. Delvenne, and J. C. Doyle, Linear-quadratic-gaussian heat engines, Proc IEEE Conference on Decision and Control, New Orleans, Louisiana, Dec. 2007

T. Cui, L. Chen, T. Ho, S. Low and L. Andrew, Opportunistic Source Coding for Data Gathering in Wireless Sensor Networks. Proceedings of IEEE MASS, October 2007.

L. Chen, S. H. Low and J. C. Doyle, Contention Control: A Game-theoretic Approach, Proc. IEEE Conference on Decision and Control, Dec 2007.

L. Chen, T. Cui, S. H. Low and J. C. Doyle, A Game-Theoretic Model for Medium Access Control, Proc. International Wireless Internet Conference (WICON), Oct 2007.

T. Cui, L. Chen and T. Ho, Energy Efficient Opportunistic Network Coding for Wireless Networks, Proc. IEEE INFOCOM, Apr 2008.

1 General Introduction and Background

Robust yet fragile (RYF) complexity is ubiquitous, appearing on all scales, from tiniest microbes and cellular subsystems to global ecosystems, to human social and technical systems, and from the early history of life's evolution through human evolution to our latest technologies and their impact on our ecosystems. Network architecture is essential to understanding RYF complexity. Protocol-based architectures (PBAs) of bio and tech networks facilitate elaborate control systems to coordinate and integrate diverse function and create coherent and global adaptation to large perturbations in their components and environments on a vast range of time and space scales, all despite implementation mechanisms that are largely decentralized and asynchronous.

Biological systems are robust and evolvable in the face of large changes in environment and components, yet can be extremely fragile to small perturbations [32]. The amazing evolution of microbes into humans (robustness of lineages on long timescales) has been punctuated by mass extinctions (extreme fragility). Diabetes, obesity, cancer, auto-immune diseases, and similar conditions are side-effects of biological control mechanisms so robust as to normally go unnoticed. Our robustly regulated and efficient metabolism that evolved when life was physical challenging and food was often scarce leads in a modern lifestyle to obesity and diabetes. Our robustly controlled physiology creates an ideal ecosystem for microbes, most of which are beneficial, but parasites can hijack our machinery for their own purposes. Our immune system works with beneficial microbes to manage this ecosystem to mutual benefit and prevents most infections, but can cause devastating autoimmune diseases, including a type of diabetes. Our complex physiology requires robust development and regenerative capacity in the adult, but this very robustness at the cellular level is turned against us in cancer.

The complexity of technology is also exploding around us, but like biology, the typical robust, simple behavior hides enormous complexity and, unfortunately, fragility. This complexity yields both benefits and catastrophes, even horrors, previously unimaginable. Organized complex societies both protect us and facilitate wars, famines, epidemics, and destruction of our ecosystems. The universal RYF nature of complex systems is well-known to domain experts such as physicians, ecologists, and systems engineers, but has been systematically studied in any unified way only recently. Indeed, much of science identifies "complexity" with chaos, criticality, and systems whose typical behavior is fragile, just the opposite of RYF.

1.1 Summary of research results

RYF complexity, high variability data, and power laws. The well-known statistical invariance of Gaussians and exponentials for low variability data make their ubiquity unsurprising. But power laws are "more normal than Normal" with even richer statistical invariants [45], and their ubiquity in the abundant high variability of architecture-based RYF systems should be equally unsurprising. Unfortunately, errors persist, with "scale free networks" (SFNs) as a recent popular example. The most (ISI) cited paper on the Internet claims that the router topology is scale-free, and thus a few 'hubs" (high connectivity routers) crucial to the overall connectivity of the system represent critical vulnerabilities (an unrecognized "Achilles heel") if attacked. Yet our work shows that SFNs are completely incompatible with both the real Internet and the underlying technology [44, 29, 45]. We showed that simple optimization-based models were consistent with data and engineer's understanding, and that the SFNs were a fiction of analysis error. The deep research challenge now is not power laws per se, but to develop rigorous tools for high variability data, particularly to distinguish the healthy role that high variability plays throughout PBA systems from their presence as symptoms of possibly excessive fragility.

HOT and a theory of RYF systems. Under the general umbrella of Highly Optimized/Organized Tolerance/Tradeoffs (HOT), we have studied a variety of optimization and architecture-based mechanisms for power laws and high variability [30, 9] as well as a variety of domain applications in wildfire ecology [50], metabolism [61], protein interaction networks, web layout, in addition to Internet router topology [29, 45]. These models are generally extreme abstractions aimed at providing both simpler and more accurate models than offered by popular models from the "chaocritical complexity" community using SFNs and SOC (Self-organized Criticality).

Architecture of biology and ecology. Progress in systems and synthetic biology combined with detailed descriptions of many components of biological networks, and the growing attention to systems biology, means the organizational principles of large-scale, complex biological networks are becoming increasingly apparent [19, 20, 31, 34, 60, 61, 42], including universal patterns of organization and architecture. While robustness (resilience) and evolvability (e.g. sorting) in ecosystems have been important topics for some time, we and our colleagues have recently published a series of studies that model ecosystem mechanisms at several levels of detail [9, 50] and explore RYF principles that relate directly to those above that occur at the cell and organism level [42].

Layered control systems in biology. Noncoding RNAs are increasingly recognized as playing a significant, and perhaps central, role in the complexity of biological control systems. In support of this, recent studies indicate that noncoding DNA, and not gene number, correlates with organism complexity. The diversity and prevalence of gene regulatory mechanisms performed by noncoding RNAs, including RNA interference, alternative splicing, and riboswitch-mediated regulation [68, 2], hints at an organizing principle guiding layered control systems in biology that balances energy, efficiency, component performance, and scaling. Understanding this and the role of layering in biology (discussed below) in the context of layered PBAs will be a major focus of our research in the future.

Architecture in technology. Advanced technology provides engineering examples of networks with complexity approaching that of biology. While the components differ from biology, we have found striking convergence at the network level of architecture and the role of layering, protocols, and feedback control in structuring complex multiscale modularity [9, 19, 44, 20, 29, 14]. The next step in this research is to deepen our comparisons within biology and ecology and between these and tech networks, to clearly explain the role of layering and protocols.

Unifying theory of complex network architecture. We believe that the apparent network-level evolutionary convergence both within biology and between biology and technology is not accidental, but follows necessarily from the requirements that both biology and technology be efficient, adaptive, evolvable, and robust to perturbations in their environment and component parts [9, 52, 18, 45, 35, 46, 61, 50]. This theory builds on and integrates decades of research in pure and applied mathematics with engineering, and specifically with theories from robust control, communications, and computational complexity, but is just the beginnings of what is necessary to fully explain the striking parallels we find in bio and tech networks, and to help design their next generation. Some specific elements are sketched below.

Unifying existing theory. The most relevant existing theoretical frameworks from thermodynamics, information theory, control theory, and computational complexity (classically associated with names such as Carnot, Shannon, Bode, and Turing) all assume incompatible architectures a priori, and thus cannot serve as a foundation to explain architectural principles. We will radically change this, building on our recent progress in unifying control theory with information theory [46] and thermodynamics. We also have several results unifying control, computational complexity, and numerical analysis theories. Our emphasis on regulatory architecture means that control theory is the natural starting point, and these results are encouraging that a complete unification is possible.

Theory of layering. The layered architecture is one of the most fundamental and influential structure of the Internet. Each layer in the protocol stack hides the complexity of the layer below and provides a service to the layer above. While the general principle of layering is widely recognized as one of the key reasons for the enormous success of the Internet, there was little quantitative understanding on a systematic process of designing layered protocol stack. We have proposed to integrate the various protocol layers into a single coherent theory, by regarding them as carrying out an asynchronous distributed computation over the network to implicitly solve a certain global optimization problem [18]. Different layers iterate on different subsets of the decision variables using local information to achieve individual optimality. Taken together, these local algorithms attempt to achieve a global objective. Such a theory will expose the interconnection between protocol layers and can be used to study rigorously the performance tradeoff in protocol layering, as different ways to distribute a centralized computation. The global behavior of protocols and networks can then be understood, and designed, through the underlying global optimization problem whose vertical and horizontal decomposition is implemented in the protocol stacks that now interact over the network.

Design and deployment of new networking protocols. This optimization-based perspective has led to the development of new TCP congestion control algorithms, FAST, that has been used by scientists to break world records of data transfers and are being further tested and deployed [63, 65]. The major existing collaboration in our team is between Doyle and Low on TCP and some related extensions.

Robustness analysis of complex systems and comparing models with data. Simulations and local sensitivity are useful and relatively cheap, and thus will remain important, but they are often insufficient, particularly when studying RYF systems. Formal tools are needed to systematically search for perturbations causing rare events or prove they will never happen, and to use data to fit parameters or invalidate models. Recent success stories can be roughly organized into three distinct axes: (1) Continuous dynamical systems, (2)Finite state systems, and (3) Networks of dynamical systems. We are working to integrate these separate and fragmented analysis techniques so as to formally reason about the global properties of systems closer to the complexity observed in bio and tech networks.

For finite state systems, computer science has developed formal logics for rigorous specification and reasoning and increasingly efficient algorithms (model checkers) to refute or verify these logical statements. For reasoning about heterogeneous networks of complex nonlinear hybrid dynamical systems, these methods are incomplete. We have recently developed complementary tools for semi-automated analysis of dynamical systems using sum-of-squares (SOS) methods (see SOSTOOLS website). Moreover, the extreme flexibility of the SOS methodology allows for extensions of these ideas to hybrid and stochastic systems [35]. In analyzing networks whose components are dynamical systems (such as in biology or the internet), network topology is uncertain and changing, and may be so large that even polynomial time algorithms for analysis are unacceptable. To avoid these problems, we seek proof methods that can hold for *arbitrary* networks by exploiting systems on graphs with special structures. Some of the more notable works are related to biological networks [61], internet topology [29, 45], stability in TCP/AQM [52, 14], flocking disturbance ecology [50] and coupled oscillators.

1.2 Architecture, Robustness, and Evolvability

Classically, architecture focused on buildings, and their function, construction and deployment, robustness to perturbations from the environment, including possibly attack, and aesthetics. Here, the term architecture focuses generally on the elements of structure and organization that are most universal, high-level, and persistent. Yet system architecture still must satisfy the classical notions of function and construction, as well as robustness and evolvability to uncertainty and change in components, function, and environment. We propose that these features provide a promising starting point in studying the architecture of complex networks in biology and technology.

The Internet hourglass terminology arose because a vast diversity of applications and hardware sit above and below a thin a thin, hidden "waist" of shared feedback control (TCP/IP). Such hourglass layering of control is a universal architectural structure. Roughly, IP controls the routes for packet flows and thus, available bandwidth. Applications split files into packets, and the TCP controls their rates and guarantees delivery. This allows "plug-and-play" between modules that obey shared protocols; any set of applications that "talks" TCP can run transparently and robustly on any set of hardware that "talks" IP.

Bowties can help visualize how protocols are organized within layers, and involve large fan-ins and -outs of energy, materials, and information via thin "knots" of universal protocols specifying carriers, building blocks, or interfaces[20]. The IP packet is a familiar example, but the application and physical layers also have their own standard interfaces, all most effective when thin knots connect diverse ends. Other examples include energy sources connected to diverse consumers via standard carriers (e.g. 110V 60Hz AC, gasoline, batteries) and socket-plug interfaces, sellers to buyers via money, and raw materials to assemblies via standardized building blocks in advance manufacturing. The currencies, carriers, intermediates, precursors, plugs, packets, conserved residues, and interfaces in the bowtie core are highly constrained by protocols. Yet their shared universality allows diverse and robust edges to adapt and evolve, as long as they have appropriate (and typically hidden) layers of feedback control.

The microbial biosphere also has a shared hourglass protocol stack for control with bowties for flow within each layer. A bacterial cell's "application layer" consists of biochemical fluxes with various bowtie protocols such as core metabolism (nutrients to biosynthesis and energy supply via precursors and carrier-based group transfers) and signal transduction (receptors to downstream effectors via conserved residue pairs). The enormous diversity of this application layer is driven by the extreme variety of environments that are experienced by the bacterial biosphere. All are supported by a shared architecture with layers and protocols that blend elements familiar from engineered architectures, but with very different hardware components (e.g. macromolecules) and more extensive and sophisticated use of feedback control.

The microbial regulatory hourglass is also layered with the lowest, physical "genomic" layer of DNA, with a control layer for DNA replication. Between the diverse application and genomic layers is a relatively thin and multiply layered waist of highly conserved core regulatory processes, also with multiple bowties. The next layer of transcriptional control regulates RNA levels, and a further layer has diverse mechanisms to control RNA action, including translation of mRNA to regulate protein levels. The bowties in each layer involve the universal, central protocols that replicate and transcribe DNA to RNA and translate to proteins via highly conserved RNA intermediates and polymerases. Most cells have an extremely complex additional layer of fast-acting control strategies to deal with fluctuating environments. Allostery, a huge suite of post-translational modifications, and the rapid changes in location of macromolecular modules enable adaptive responses at the application layer to environmental signals or alterations on rapid time scales. Note that mature red blood cells have only this upper layer, reinforcing the biological relevance of this layering abstraction.

These universally shared sets of protocols are more fundamental and invariant than the modules whose control and evolution they facilitate, and genes that "talk" these central protocols can move by horizontal gene transfer (HGT), accelerating evolution in a kind of "bacterial internet." Like the technological Internet, the functionality of the resulting new proteins that result is enhanced by having additional shared application layer protocols, such as group transfers and carriers in metabolism, and conserved residue pairs in signal transduction. Thus selection acting at the protocol level could evolve and preserve shared architecture, essentially "evolving evolvability." While HGT is perhaps the most dramatic mechanism by which bacteria can "facilitate variation" through large but functional rearrangements of their genome, there are a host of other methods and more are being discovered, particularly at the RNA level.

On longer time scales within and across generations, the sequence of the DNA itself can change, not only through random point mutation, but also through highly structured mechanisms that facilitate the generation of large, adaptive diversity. Furthermore, as biologists dig deeper past the superficiality of sequence data into the complexity of regulation, they unearth additional layers of control that are fundamentally similar to those in advanced technologies, with the explosion of interest in RNA-based regulation particularly striking.

Evolvability, from microbes to ecosystems to IP-based networks illustrate how dramatic, novel, dynamic, distributed changes on all scales of time and space can also be coherent, responsive, functional and adaptive. New genes and pathways, laptops and applications, even whole networks, can all "plug-and-play", as long as they obey the appropriate protocols. All of life and advanced technologies rely on protocol-based architecture, and evolution is also robustness, but to lineages of systems or organisms, on long time scales, and often to very large uncertainties. While we claim our emphasis on architecture is totally consistent with the latest and deepest thinking in biology, potentially what is most radical and controversial about our approach are the implications for evolution, and for the construction of synthetic systems.

The "fitness landscape" has long been emphasized as a central concept in mainstream theories of evolution, involving small random mutation but deterministic selection. That is, repeated small, random, and typically deleterious genotype variation yield also small phenotype variation which is then acted on by an essentially deterministic fitness selection process to yield incrementally higher fitness organisms, hence the image of climbing on a "fitness landscape." The emphasis is on the accumulation of accident and drift towards higher fitness. This view of evolution as tinkerer is conceptually elegant, mathematically tractable, and connects naturally with molecular biology's traditional tool of genetic screens in idealized laboratory environments. Yet it plays a small role in the evolution of advanced technologies, the hype surrounding genetic algorithms notwithstanding, and may also actually play only a small background role in evolution in natural environments.

PBAs accelerate evolution in myriad ways, but many examples are now familiar and illustrative. It is well-known that the rapid acquisition of bacterial antibiotic resistance does not occur by the incremental accumulation of point mutations, but occurs via large-scale rearrangements, including HGT. This would not be possible without the modularity that PBAs facilitate. Thus variation need not be small and slow, but can be fast, large, yet nevertheless functional. Another familiar example is the rapid evolution and diversification of dogs from wolves, involving not the gradual accumulation of a large number of independent mutations, but a few small changes in regulatory processes that yield large changes in phenotype via coordinated developmental changes. Thus with PBAs, small genotype change can yield large but functional phenotype change.

In natural environments fitness and selection is highly stochastic and accidental, and protocol-based architectures "facilitate variation" that while random in origin, can be large, structured, and most importantly, highly likely to be adaptive. This makes significantly different predictions that are more consistent with observations than conventional theories. It also connects more directly with the evolvability of technology. This use of conserved core processes with small regulatory changes to yield large phenotype changes is now widely believed to be responsible for the large diversity in plant and animal life despite largely shared protein-coding regions in their genomes. Evolvability is RYF as well. That is, the evolvability facilitated by PBAs requires the adoption of universal protocols that are subsequently difficult or impossible to change (captured elegantly by Gerhart and Kirschner's "constraints that deconstrain" and "facilitated variation").

2 Selected technical details

2.1 Layering and optimization

Communication networks are complex systems consisting of intelligent units such as PCs that have computing and communicating capabilities. The functioning of the network as a whole is made possible by all kinds of protocols that integrate these individual units together. As networks adopt a layered structure, protocols for different layers are usually optimized and implemented separately, and then interconnected, often in ad hoc manner. Though justified by the success of many communication networks, this "layered" design methodology does not perform well for wireless networks because of time-varying channel, contention based channel access and mobility. In order for wireless networks to provide better performance, we must re-think the protocol stack as a whole, and exploit the interactions among various layers to do cross-layer design.

The approach of protocol as distributed solution to some global optimization problem through dual decomposition has been successfully applied to TCP congestion control. The key innovation in this line of work is to view the network as an optimization solver and congestion control protocol as a distributed algorithm solving a Network Utility Maximization (NUM) problem. This approach has recently been substantially extended from an analytic tool of reverse-engineering TCP congestion control to a general approach to understand interactions across layers. Application needs form the objective function, i.e., network utility to be maximized, and the restrictions in the communication infrastructure are translated into many constraints of a generalized network utility maximization problem. Such problems in general may be very difficult nonlinear, nonconvex optimization with integer constraints. There are many different ways to decompose a given problem, each of which corresponds to a different layering scheme. These decomposition (i.e., layering) schemes have different trade-offs in efficiency, robustness, and asymmetry of information and control, thus some are "better" than others depending on the criteria set by the network users and operators.

Wireless networks We apply this approach to design an overall framework for the protocol architecture in ad hoc wireless networks, with the goal of achieving efficient resource allocation through cross-layer design [13, 12]. Our current theory integrates three functions - congestion control, routing, and scheduling - in transport, network, and link layers into a coherent framework and makes transparent their interactions not only vertically across the protocol stack, but also horizontally across multiple network nodes. These three functions interact through and are regulated by congestion prices so as to achieve global optimality, even in a time-varying environment. The structural simplicity of the underlying optimization problem also leads to simple and robust equilibrium and dynamic behaviors. Even though this framework does not provide all the design and implementation details, it helps us understand issues, clarify ideas, and suggests directions leading to better and more robust designs for ad hoc wireless networks.

Wireless scheduling We also study wireless scheduling. Wireless scheduling is very challenging, due to the interdependence of wireless links and the time-varying nature of wireless channel. In [15, 16], we propose a general dual scheduling algorithm that uses rate control and queue-length based scheduling to allocate resources for a generalized switch with interdependent and time-varying service capabilities. We first consider a saturated system in which each user has infinite amount of data to be served. We prove the asymptotic optimality of the dual scheduling algorithm for such a system, which says that the vector of average service rates of the scheduling algorithm maximizes some aggregate concave utility functions. As the fairness objectives can be achieved by appropriately choosing utility functions, the asymptotic optimality establishes the fairness properties of the dual scheduling algorithm. We next consider a system with exogenous arrivals, i.e., data flows of finite size arrive at the system randomly. For such a system, we propose a modified dual scheduling algorithm that stabilizes the system whenever the input rates are within the feasible rate region and is then throughput-optimal, i.e., achieves 100% throughput. The dual scheduling algorithm motivates a new architecture for scheduling, in which an additional queue is introduced to interface the user data queue and the time-varying server and to modulate the scheduling process, so as to achieve different performance objectives. This new scheduling algorithm and architecture is going to find many applications from fair scheduling in cellular networks to Quality of Service scheduling in packet switch in future wireless internet.

A game-theoretic approach to contention control We developed a game-theoretic model, called random access game, which provides a unique perspective to understand existing contention based medium access control protocols and a general framework to guide the design of new ones to improve the system performance: medium access protocol can be interpreted as and designed according to distributed strategy update algorithm achieving the equilibrium of random access game [17, 10]. The random access game is a rather general construction, as it can be reverse-engineered from existing MAC protocols, forward-engineered from desired operating points, or designed based on heuristics [10]. Medium access methods derived from concrete random access game models achieve superior performance over the standard IEEE 802.11 DCF and can provide flexible service differentiations [17].

Practical networking with network coding Recent advances in network coding have shown great potential for efficient information multicasting in communication networks, in terms of both network throughput and network management. We have proposed optimization based models to study the important issue of integrating network coding with congestion control [11], wireless scheduling [21] and multipath routing [22]. In [11], we address the problem of rate control at end-systems for network coding based multicast flows. We develop two adaptive rate control algorithms for wired networks with given coding subgraphs and without given coding subgraphs, respectively. With random network coding, both algorithms can be implemented in a distributed manner, and work at transport layer to adjust source rates and at network layer to carry out network coding. We also studied inter-session network coding. In [23], we proposed a general session decomposition approach for inter-session network coding, which transforms an inter-session network coding problem into an equivalent intra-session network coding problem, and proposed an optimization based model to guide the design of an oblivious backpressure algorithm of dynamic scheduling and inter-session coding for wireless networks with multiple unicast flows. A practical protocol was designed to implement the oblivious backpressure algorithm, and numerical experiments showed it largely reduces power consumption over all existing algorithms, by exploiting multiple-reception gain and using network coding.

Distributed data gathering Correlated data gathering is usually formulated as a minimum cost problem (e.g., to minimize the energy usage). Many works are based on the use of distributed source coding such as Slepian-Wolf coding. There are two main problems with these works. First, they require coordination among the sources to guarantee that the source rates lie in the Slepian-Wolf region and hence do not admit distributed implementation. Second, for the multicast case where the data need to be sent to multiple sinks, they involve finding a minimum cost Steiner tree, which is known to be NP-complete. In [25], we propose fully distributed algorithms for lossless data gathering for multiple multicast sessions with correlated sources, in which the sinks control transmission rates across the sources via local updates that propagate back to the sources. Network coding is also used to improve the transport efficiency and to avoid the hard problem of finding minimum cost Steiner tree. In [26], by formulating it as a distortion-utility optimization problem, we propose a distributed algorithm for lossy data gathering for correlated sources. The resulting receiver-driven algorithm adjusts distortion levels according to the feedback from the sinks, and hence does not require the coordination among the sources.

Wireless sensor networks Another research focus is on wireless sensor networks. Due to limited energy and bandwidth, one of the most important issues in wireless sensor networks is how to achieve efficient data gathering. Previous data gathering schemes did not exploit wireless broadcast advantage, and thus resulted in large efficiency loss. In [24], we proposed a novel opportunistic source coding and opportunistic routing (OSCOR) protocol for correlated data gathering in wireless sensor networks. This new scheme improves data gathering efficiency by exploiting opportunistic data compression and cooperative diversity with wireless broadcast advantage. The design of OSCOR involves several challenging issues across different network protocol layers. At MAC layer, sensor nodes need to coordinate wireless transmission and packet forwarding to exploit multiuser diversity in packet reception. At network layer, in order to achieve high diversity and compression gains, routing must be based on a metric that is dependent on not only link-quality but also compression opportunities. At application layer, sensor nodes need a distributed source coding algorithm that has low coordination overhead and does not require the source distributions to be known. OSCOR provides practical solutions to these challenges incorporating a slightly modified 802.11 MAC, a distributed source coding scheme based on Lempel-Ziv code and network coding, and a node compression ratio dependent metric combined with a modified Dijkstra's algorithm for path selection. We evaluate the performance of OSCOR through simulations, and show that OSCOR reduces the number of transmissions by nearly 25% compared with existing greedy scheme in small networks. We expect a large gain in large networks.

2.2 Topology of complex networks

During the last decade, significant efforts have been made toward improving our understanding of the topological structures underlying complex networks and illuminating some of the intriguing large-scale properties exhibited by these systems. The dominant theme of these efforts has been on studying the graph-theoretic properties of the corresponding connectivity structures and on developing universal theories and models that transcend systemspecific details and describe the different systems well in a statistical sense.

Towards a Theory of Scale-Free Graphs There is a large, popular, and growing literature on "scale-free" networks with the Internet along with metabolic networks representing perhaps the canonical examples. While this has in many ways reinvigorated graph theory, there is unfortunately no consistent, precise definition of scale-free graphs and few rigorous proofs of many of their claimed properties. In fact, it is easily shown that the existing theory has many inherent contradictions and that the most celebrated claims regarding the Internet and biology are verifiably false. In this paper, we introduce a structural metric that allows us to differentiate between all simple, connected graphs having an identical degree sequence, which is of particular interest when that sequence satisfies a power law relationship. We demonstrate that the proposed structural metric yields considerable insight into the claimed properties of SF graphs and provides one possible measure of the extent to which a graph is scale-free. This structural view can be related to previously studied graph properties such as the various notions of self-similarity, likelihood, betweenness and assortativity. Our approach clarifies much of the confusion surrounding the sensational qualitative claims in the current literature, and offers a rigorous and quantitative alternative, while suggesting the potential for a rich and interesting theory.

Diversity of graphs with highly variable connectivity A popular approach for describing the structure of many complex networks focuses on graph theoretic properties that characterize their large-scale connectivity. While it is generally recognized that such descriptions based on aggregate statistics do not uniquely characterize a particular graph and also that many such statistical features are interdependent, the relationship between competing descriptions is not entirely understood. This paper lends perspective on this problem by showing how the degree sequence and other constraints e.g., connectedness, no self-loops or parallel edge on a particular graph play a primary role in dictating many features, including its correlation structure. Building on recent work, we show how a simple structural metric characterizes key differences between graphs having the same degree sequence. More broadly, we show how often implicit choice of a background set against which to measure graph features has serious implications for the interpretation and comparability of graph theoretic descriptions.

2.3 Physical hard limits in networked control systems

Fundamental limitations and hard limits are at the core of many branches of engineering and science. These limits provide bounds on what we can and cannot achieve. Examples can be found in estimation theory, information theory, and control theory. Technology from all of these branches of engineering are used together and in parallel in modern networked control systems. Currently a lot of research effort is put into combining the hard limits results of information and control theory. In the physical sciences there are also well-known hard limits. We call these *physical hard limits*. We argue that since all components in a networked system need to be implemented by physical devices, the physical hard limits give hard limits on the performance of the entire networked system.

We focus on physical hard limits arising from thermal noise and from devices with finite available energy in this project. The specific goal is to quantify hard limits of physical implementations of typical devices such as actuators, measurement devices, controllers, and to interconnections of them. Our preliminary results indicate that a good approach to quantify these limitations is to match lossless dynamical systems to device models. As an application for the results we plan to use a control-theoretic heat engine that transfers uncertain energy (heat) into useful energy (work) using feedback control. That is an interesting problem in its own right, with close connections to thermodynamics and Maxwell's demon.

In this section we discuss our attempt to bring in some of the physical limitations into the picture of hard limits for networked control systems. Since current research has mostly been focused on limits arising from information theory and control theory, this should help broaden our understanding of networked control systems.

Background and motivation In estimation theory, the Cramér-Rao inequality gives a lower bound on the covariance of the estimation error, in information theory a theorem by Shannon states that channel capacity provides an upper limit on the communication rate, and in control theory Bode's integral puts bounds on the sensitivity function. These are just some examples of fundamental limitations and hard limits relevant to an engineer working with networked control systems. For an overview of these bounds, see [58]. Current research has put a lot of effort into generalizing and combining various fundamental limitations, see [51, 62, 56, 47], and references therein. In the physical sciences there are also some well-known hard limits. For the purpose of our work, the fluctuation-dissipation theorem [41] and Carnot's theorem [64] seem the most relevant ones. These are reviewed next.

Consider the standard model of a linear electrical resistor of resistance $R > 0$,

$$v = Ri, \tag{1}$$

where i is the current and v the voltage drop. This is a dissipative model [66, 67] since the power input vi is always non-negative. The fluctuation-dissipation theorem says that all dissipative devices exhibit noise, and a special instance is the Johnson-Nyquist noise of a resistor [64]:

$$v = Ri + \sqrt{2k_BTR}m(t), \tag{2}$$

where T is the absolute temperature of the resistor, k_B Boltzmann's constant, and $m(t)$ is unit-intensity white noise. The model (2) captures the empirical fact that the voltage across a resistor is noisy. Hence, whenever a resistor is used in a networked system, there will be a physical noise source in the system. This is an example of a physical hard limit: No dissipation without noise. The model (2) is well known. It is less clear what the physical hard limits are for active devices, such as a *negative resistor* ($R < 0$). Since standard components in a control system, such as actuators, controllers, and measurement devices often are active, we want to understand what the hard limits are for these.

Another physical hard limit is found in thermodynamics. There heat engines are defined as devices that convert uncertain energy (heat) into directly useful energy (work). Using the laws of thermodynamics, one can derive an upper bound on the efficiency and construct a heat engine, the Carnot heat engine, that achieves this bound. The bound is simple, $1 - T_{\text{cold}}/T_{\text{hot}}$, where T_{cold} and T_{hot} are the temperatures of the heat sources the engine can exchange heat with, see, for example [64]. The Carnot heat engine is a theoretical construct that operates infinitely slowly and in quasi equilibrium. Hence, the basic theory does not answer how well a heat engine can do over finite time intervals. To study this, we have to put more assumptions on the environment and the engine. For a control engineer it is natural to study these problems in a dynamical systems setting. Some control-theoretic approaches to thermodynamics and analysis of heat engines are found in [4, 40, 27]. A control-theoretic treatment of the Carnot heat engine that uses the noisy resistor model (2) is given in [4]: Consider the switched capacitor-resistor circuit modeled by

$$\frac{d}{dt}[C(t)v(t)] = -g_i v(t) + \sqrt{2k_BT_ig_i}m(t), \quad g_i \in \{1/R_1, 1/R_2, 0\}, \quad T_1 \geq T_2, \tag{3}$$

and mechanical work can be extracted by varying the distance between the capacitor plates. Then the capacitance $C(t)$ is time varying and the expected work rate is

$$w_{mech}(t) = \frac{1}{2}\dot{C}(t)\mathbf{E}v(t)^2.$$

By open-loop variations of the capacitance $C(t)$ and operating the switch g_i properly in open loop between the warm resistor R_1 of temperature T_1 and the cold resistor R_2 of temperature T_2, we can extract work from heat. It can be shown, see [4, 27], that the efficiency bound

$$\frac{\int_{cycle} w_{mech} dt}{\int_{cycle} q_1 dt} \leq 1 - \frac{T_2}{T_1}, \tag{4}$$

holds for all cycles. Here q_1 is the heat flow from the warmer resistor. For infinitely long cycles, we achieve equality in (4) by essentially using the operating scheme of the Carnot heat engine. This is a physical hard limit: It is physically impossible to achieve higher efficiency using open-loop control. What this model does not capture is the possibility of using measurements and feedback for increasing the efficiency. A feedback controller could here act as a Maxwell's demon, see, for example [49]. Maxwell's demon is a fictional controller that makes the system violate the second law of thermodynamics ("total entropy always increases"). The usual solution to this problem is to claim that a real physical implementation of a Maxwell's demon always leads to heat generation, for example by referring to Landauer's principle [43], and that will save the second law from violation. See also the discussion in [27]. However, even for a concrete and simple model such as (3), it is not clear how to quantitatively and rigorously prove such a statement given a controller. Our hopes are that a theory for physical hard limits on implementation of physical devices will ultimately help to make such an analysis possible.

The Statistical Mechanics of Fluctuation-Dissipation and Measurement Back Action [57] In this paper, we take a control-theoretic approach to answering some standard questions in statistical mechanics. A central problem is the relation between systems which appear macroscopically dissipative but are microscopically lossless. We show that a linear macroscopic system is dissipative if and only if it can be approximated by a linear lossless microscopic system, over arbitrarily long time intervals. As a by-product, we obtain mechanisms explaining Johnson-Nyquist noise as initial uncertainty in the lossless state, as well as measurement back action and a trade-off between process and measurement noise.

Thermodynamics of Linear Systems [27] In this paper, we rigorously derive the main results of thermodynamics, including Carnot's theorem, in the framework of time-varying linear systems.

Linear-Quadratic-Gaussian Heat Engines We also study the problem of extracting work from heat flows. In thermodynamics, a device doing this is called a heat engine. A fundamental problem is to derive hard limits on the efficiency of heat engines. Here we construct a linear-quadratic-Gaussian optimal controller that estimates the states of a heated lossless system. The measurements cool the system, and the surplus energy can be extracted as work by the controller. Hence, the controller acts like a Maxwell's demon. We compute the efficiency of the controller over finite and infinite time intervals, and since the controller is optimal, this yields hard limits. Over infinite time horizons, the controller has the same efficiency as a Carnot heat engine, and thereby it respects the second law of thermodynamics. As illustration we use an electric circuit where an ideal current source extracts energy from resistors with Johnson-Nyquist noise.

Goals and Challenges In [57]) it is shown that lossless approximation is a useful tool for quantifying hard limits. One goal of our work is to

a) develop methodology for linear and nonlinear lossless model approximation of dissipative and active models.

The plan is to continue using control-theoretic methods. Once we have a better understanding of lossless approximation, the next step is to

b) find physical hard limits, arising from thermal noise and finite available energy, for implementations of simple measurement devices, actuators, and controller models.

Next it is natural to study interconnections and networks of such devices. Hence, we plan to

c) quantify hard limits for interconnections and networks of lossless models.

An example of such a hard limit is given in the preliminary work in [57]: A lossless implementation of a dissipative linear measurement device results in a back action effect on the measured lossless system. As an application we also plan to study a control-theoretic heat engine. In particular, it will be interesting to see how efficiency bounds like (4) are effected by feedback control when all devices are assumed to be lossless, exhibiting noise and/or having only finite amount of energy. As already discussed, it is often argued that it is heat losses in implementations of controllers that prevent the second law of thermodynamics from being violated in such cases. It would be interesting to be able to give such statements a rigorous control-theoretic meaning, even if only for a simple circuit.

Finally, this project has close connections to many other scientific fields, and we have to

d) closely study related work in statistical mechanics, information theory, quantum mechanics, and circuit theory.

It is also likely that a behavioral model setting [54] is more natural for treating the interconnections of models.

2.4 From Sparsity and Compressed Sensing to Matrix Rank Minimization

Low-rank matrices and sparse vectors In many applications, notions such as order, complexity, or dimension of a model or design can be expressed by means of the rank of an appropriate matrix. For example, a low-rank matrix could correspond to a low-degree statistical model for a random process (e.g., factor analysis), a low-order controller for a plant [33], a low-order realization of a linear system [37], or a low-dimensional embedding of data in Euclidean space [38]. If the set of feasible models or designs is affine in the matrix variable, choosing the simplest model can be cast as an (affine) *rank minimization problem*, which can be written as

$$\begin{array}{ll} \text{minimize} & \operatorname{rank}(X) \\ \text{subject to} & \mathcal{A}(X) = b, \end{array} \tag{5}$$

where $X \in \mathbb{R}^{m\times n}$ is the decision variable, and the linear map $\mathcal{A} : \mathbb{R}^{m\times n} \to \mathbb{R}^p$ and vector $b \in \mathbb{R}^p$ are given. In general, problem (5) is a challenging nonconvex optimization problem for which no polynomial time solution procedure is known.

As a special case, when the matrix variable is constrained to be diagonal, problem (5) reduces to finding the *sparsest vector* in an affine subspace. This problem is also referred to as *cardinality minimization*, since we seek the vector whose support has the smallest cardinality. Note that while for the sparsest vector problem there is in principle an (admittedly inefficient) way of producing candidate solutions by enumerating finitely many supports, for rank minimization no such finite procedure exists.

In certain instances with very special structure, the rank minimization problem can be solved by using the singular value decomposition, or can be exactly reduced to the solution of linear systems [48, 53]. For the general case, several heuristic approaches exist. One group of heuristic methods are based on local optimization, and include methods such as alternating projections [39] and alternating LMIs [59]. Another approach is to replace the rank objective with a convex function and solve the resulting convex problem, which can be done efficiently. One such method is the *nuclear norm* heuristic [36], which minimizes the nuclear norm, or the sum of the singular values of the matrix, instead of its rank. The nuclear norm has been shown to be the convex envelope of the rank function over the unit ball (matrices with norm less than one), which gives an interpretation and theoretical justification for the heuristic. When the matrix variable is symmetric and positive semidefinite, this is the same as the trace heuristic used often in the control community (e.g., [3]). The nuclear norm heuristic has been observed to produce very low-rank (and sometimes optimal) solutions in practice, but a theoretical characterization of when it is exact has not been previously available.

In the case when the matrix is diagonal, the sum of the singular values is equal to the sum of the absolute values (i.e., the ℓ_1 norm) of the diagonal elements. Minimization of the ℓ_1 norm is a well-known heuristic for the cardinality minimization problem . This is a convex optimization problem, that can be reformulated, for instance, in terms of a standard linear programming problem. There is much practical evidence that, at least for certain classes of problems, the ℓ_1 heuristic yields high-quality (or even optimal) sparse solutions.

Connection with compressed sensing In recent years, there has been much research activity centered on the characterization of instances for which the ℓ_1 heuristic can be *a priori* guaranteed to yield the optimal solution. These remarkable results ensure that what in principle looks like a very difficult combinatorial optimization problem, can nevertheless be solved by efficient convex optimization techniques, in polynomial time. These techniques provide the foundations of the recently developed *compressed sensing* [28] or *compressive* *sampling* [5] frameworks for measurement, coding, and signal estimation. As has been shown by a number of research groups (e.g., [8, 6, 7, 1]), the ℓ_1 heuristic for sparsity minimization provably works whenever the sensing matrix has certain "basis incoherence" properties, and in particular, when it is randomly chosen according to certain specific ensembles.

The fact that the ℓ_1 heuristic is a special case of the nuclear norm heuristic suggests that these results from the compressed sensing literature might be extendable to provide guarantees about the nuclear norm heuristic for the rank minimization problem. In our work described below, we show that the parallels are surprisingly strong.

Guaranteed minimum-rank solutions of matrix equations via nuclear norm minimization In [55] we show that if a certain restricted isometry property holds for the linear transformation defining the constraints, the minimum rank solution can be recovered by solving a convex optimization problem, namely the minimization of the nuclear norm over the given affine space. We present several random ensembles of equations where the restricted isometry property holds with overwhelming probability, provided the codimension of the subspace is $O(r(m+n)\log mn)$, where m, n are the dimensions of the matrix, and r is its rank.

The techniques used in our analysis have strong parallels in the compressed sensing framework. We discuss how affine rank minimization generalizes this pre-existing concept and outline a dictionary relating concepts from cardinality minimization to those from rank minimization. We also discuss several algorithmic approaches to solving the norm minimization relaxations, and illustrate our results with numerical examples.

References

1. R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Preprint*, 2007. `http://dsp.rice.edu/cs/jlcs-v03.pdf`.

2. T. S. Bayer and C. D. Smolke. Programmable ligand-controlled riboregulators of eukaryotic gene expression. *Nature Biotechnology*, 23(3):337–343, February 2005.

3. C. Beck and R. D'Andrea. Computational study and comparisons of LFT reducibility methods. In *Proceedings of the American Control Conference*, 1998.

4. R. W. Brockett and J. C. Willems. Stochastic control and the second law of thermodynamics. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1007–1011, San Diego, California, 1978.

5. E. J. Candès. Compressive sampling. In *International Congress of Mathematicians. Vol. III*, pages 1433–1452. Eur. Math. Soc., Zürich, 2006.

6. E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications of Pure and Applied Mathematics*, 59:1207–1223, 2005.

7. E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.

8. E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.

9. J. M. Carlson and J. C. Doyle. Complexity and robustness. *PNAS*, 99:2538–2545, 2002.

10. L. Chen, T. Cui, S. H. Low, and J. C. Doyle. A game-theoretic model for medium access control. In *Proc. of International Wireless Internet Conference (WICON)*, October 2007.

11. L. Chen, T. Ho, S. H. Low, M. Chiang, and J. C. Doyle. Optimization based rate control for multicast with network coding. In *Proc. of IEEE Infocom*, May 2007.

12. L. Chen, S. Low, M. Chiang, and J. Doyle. Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks. In *Proceedings of IEEE Infocom*, 2006.

13. L. Chen, S. Low, and J. Doyle. Joint congestion control and media access control design for ad hoc wireless networks. In *Proceedings of IEEE Infocom*, 2005.

14. L. Chen, S. H. Low, M. Chiang, and J. C. Doyle. Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks. In *IEEE Infocom*, 2006.

15. L. Chen, S. H. Low, and J. C. Doyle. Dual scheduling algorithm in a generalized switch: Asymptotic optimality and throughput optimality. In I. Elhanany and M. Hamdi, editors, *High-Performance Packet Switching Architectures*. Springer, August 2006.

16. L. Chen, S. H. Low, and J. C. Doyle. On asymptotic optimality of dual scheduling algorithm in a generalized switch. In *Proc. of International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*, April 2006.

17. L. Chen, S. H. Low, and J. C. Doyle. Contention control: A game-theoretic approach. In *Proc. of IEEE CDC*, December 2007.

18. M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle. Layering as optimization decomposition. *Proceedings of the IEEE*, 95, 2007.

19. M. Csete and J. Doyle. Reverse engineering of biological complexity. *Science*, 295(5560):1664–1669, 2002.

20. M. Csete and J. Doyle. Bowties, metabolism, and disease. *Trends BioTech*, 22:446–450, 2004.

21. T. Cui, L. Chen, and T. Ho. Distributed optimization in wireless networks using broadcast advantage. In *Proc. of IEEE CDC*, December 2007.

22. T. Cui, L. Chen, and T. Ho. Optimization based rate control for multicast with network coding: A multipath formulation. In *Proc. of IEEE CDC*, December 2007.

23. T. Cui, L. Chen, and T. Ho. Energy efficient opportunistic network coding for wireless networks. In *Proc. of IEEE Infocom*, April 2008.

24. T. Cui, L. Chen, T. Ho, S. Low, and L. Andrew. Opportunistic source coding for data gathering in wireless sensor networks. In *Proc. of IEEE Conf. on Mobile Ad-hoc and Sensor Systems*, October 2007.

25. T. Cui, T. Ho, and L. Chen. Distributed minimum cost multicasting with lossless source coding and network coding. In *Proc. of IEEE CDC*, December 2007.

26. T. Cui, T. Ho, and L. Chen. On distributed distortion optimization for correlated sources with network coding. In *Proc. of IEEE ISIT*, June 2007.

27. J.-C. Delvenne, H. Sandberg, and J. C. Doyle. Thermodynamics of linear systems. To appear at the European Control Conference 2007.

28. D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.

29. J. Doyle and et al. The “robust yet fragile” nature of the internet. *PNAS*, 102(41), 2005.

30. J. C. Doyle and J. M. Carlson. Power laws, Highly Optimized Tolerance and generalized source coding. *Physical Review Letters*, 84(24):5656–5659, 2000.

31. J. C. Doyle and M. Csete. Motifs, stability, and control. *PLOS Biology*, 3, 2005.

32. J. C. Doyle and M. E. Csete. Rules of engagement. *Nature*, 2007.

33. L. El Ghaoui and P. Gahinet. Rank minimization under LMI constraints: A framework for output feedback problems. In *Proceedings of the European Control Conference*, 1993.

34. H. El-Samad, H. Kurata, J. C. Doyle, C. A. Gross, and M. Khammash. Surviving heat shock: Control strategies for robustness and performance. *PNAS*, 102(8):2736–2741, 2005.

35. H. El-Samad, S. Prajna, A. Papachristodoulou, J. Doyle, and M. Khammash. Advanced methods and algorithms for biological network analysis. *Proc of the IEEE*, 90(4):832–853, 2006.

36. M. Fazel. *Matrix Rank Minimization with Applications.* PhD thesis, Stanford University, 2002.

37. M. Fazel, H. Hindi, and S. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the American Control Conference*, 2001.

38. M. Fazel, H. Hindi, and S. Boyd. Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices. In *Proceedings of the American Control Conference*, 2003.

39. K. M. Grigoriadis and E. B. Beran. Alternating projection algorithms for linear matrix inequalities problems with rank constraints. In L. El Ghaoui and S. Niculescu, editors, *Advances in Linear Matrix Inequality Methods in Control*, chapter 13, pages 251–267. SIAM, 2000.

40. W. M. Haddad, V. S. Chellaboina, and S. G. Nersesov. *Thermodynamics: A Dynamical Systems Approach*. Princeton University Press, 2005.

41. R. Kubo. The fluctuation-dissipation theorem. *Reports on Progress in Physics*, 29(1):255–284, 1966.

42. H. Kurata, H. El Samad, R. Iwasaki, H. Ohtake, and J. C. Doyle. Module-based analysis of robustness tradeoffs in the heat shock response system. *PLoS Comput Biol*, 2, 2006.

43. R. Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961.

44. L. Li, D. Alderson, J. Doyle, and W. Willinger. A first-principles approach to understanding the internet's router-level topology. *ACM SIGCOMM*, 2004.

45. L. Li, D. Alderson, J. Doyle, and W. Willinger. Towards a theory of scale-free graphs: Definition, properties, and implications. *Internet Mathematics*, 2(4), 2005.

46. N. C. Martins, M. A. Dahleh, and J. C. Doyle. Fundamental limitations of disturbance attenuation in the presence of side information. *IEEE Trans. Automatic Control*, 2007.

47. N. C. Martins, M. A. Dahleh, and J. C. Doyle. Fundamental limitations of disturbance attenuation in the presence of side information. *IEEE Transactions on Automatic Control*, 52(1):56–66, January 2007.

48. M. Mesbahi and G. P. Papavassilopoulos. On the rank minimization problem over a positive semidefinite linear matrix inequality. *IEEE Transactions on Automatic Control*, 42(2):239–243, 1997.

49. S. K. Mitter and N. J. Newton. Information and entropy flow in the Kalman-Bucy filter. *Journal of Statistical Physics*, 118:145–176, 2005.

50. M. A. Moritz, M. E. Morais, L. A. Summerell, J. M. Carlson, and J. C. Doyle. Wildfires, complexity, and highly optimized tolerance. *PNAS*, 102, 2005.

51. G. N. Nair and R. J. Evans. Stabilizability of stochastic linear systems with finite feedback data rates. *SIAM Journal on Control and Optimization*, 43(2):413–436, July 2004.

52. A. Papachristodoulou, L. Li, and J. C. Doyle. Methodological frameworks for large-scale network analysis and design. *SIGCOMM Comput. Commun. Rev.*, 34(3):7–20, 2004. 3.

53. P. A. Parrilo and S. Khatri. On cone-invariant linear matrix inequalities. *IEEE Trans. Automat. Control*, 45(8):1558–1563, 2000.

54. J. W. Polderman and J. C. Willems. *Introduction to Mathematical Systems Theory — A Behavioral Approach*. Springer, 1997.

55. B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of matrix equations via nuclear norm minimization. In *Proceedings of the 2007 Allerton Conference*, 2007. Preprint of full version available at http://lanl.arxiv.org/abs/0706.4138.

56. H. Sandberg and B. Bernhardsson. A Bode sensitivity integral for linear time-periodic systems. *IEEE Transactions on Automatic Control*, 50(12):2034–2039, Dec. 2005.

57. H. Sandberg, J.-C. Delvenne, and J. C. Doyle. The statistical mechanics of fluctuation-dissipation and measurement back action. To appear at the American Control Conference 2007. Preprint available at http://arxiv.org/abs/math.DS/0611628.

58. M. M. Seron, J. H. Braslavsky, and G. C. Goodwin. *Fundamental Limitations in Filtering and Control*. Springer, London, 1997.

59. R. E. Skelton, T. Iwasaki, and K. Grigoriadis. *A Unified Algebraic Approach to Linear Control Design*. Taylor and Francis, 1998.

60. J. Stelling, U. Sauer, Z. Szallasi, F. Doyle, and J. Doyle. Robustness of cellular functions. *Cell*, 118:675–685, 2004.

61. R. Tanaka, M. Csete, and J. Doyle. Highly optimised global organization of metabolic networks. *IEE Proceedings Systems Biology*, 152(4):179–184, 2005.

62. S. Tatikonda, A. Sahai, and S. K. Mitter. Stochastic linear control over a communication channel. *IEEE Transactions on Automatic Control*, 49(9):1549–1561, September 2004.

63. J. Wang, L. Li, S. H. Low, and J. C. Doyle. Cross-layer optimization in TCP/IP networks. *IEEE/ACM Trans. on Networking*, 13:582–595, 2005.

64. G. H. Wannier. *Statistical Physics*. Dover Publications, 1987.

65. D. X. Wei, C. Jin, S. H. Low, and S. Hegde. Fast tcp: motivation, architecture, algorithms, performance. *IEEE/ACM Transactions on Networking*, 14:1246–1259, 2006.

66. J. C. Willems. Dissipative dynamical systems part I: General theory. *Archive for Rational Mechanics and Analysis*, 45:321–351, 1972.

67. J. C. Willems. Dissipative dynamical systems part II: Linear systems with quadratic supply rates. *Archive for Rational Mechanics and Analysis*, 45:352–393, 1972.

68. M. N. Win and C. D. Smolke. From the cover: A modular and extensible rna-based gene-regulatory platform for engineering cellular behavior. *PNAS*, 104:14283–8, 2007.